Wherefore, what is claimed is:

1. A computer-implemented face recognition process for identifying a person depicted in an input image, comprising using a computer to perform the following process actions:

creating a database of a plurality of model image characterizations, each of which represents the face of a known person that it is desired to identify in the input image as well as the person's face pose;

comparing a region depicting the face of a person, which has been extracted from said input image and characterized in a manner similar to the plurality of model images, to the plurality of model image characterizations;

ascertaining which of the plurality of model image characterizations most closely matches the similarly characterized input image region; and

designating the input image region to be the person associated with the most closely matching model image characterization if a degree of similarity exceeds a prescribed threshold.

2. The process of Claim 1, further comprising a process action of specifying that the person designated as corresponding to the input image region has the face pose associated with the most closely matching model image characterization.

3. The process of Claim 1, further comprising a process action for designating the input image region to be an unknown person if the degree of similarity between the characterized input region and the most closely matching model image characterization does not exceed the prescribed threshold.

4. The process of Claim 2, wherein the comparing, ascertaining, designating and specifying process actions comprises the actions of:

training a neural network ensemble to identify the person associated with the characterized input image region and their face pose; and

employing the neural network ensemble to identify the person associated with the characterized input image region, and their face pose.

5. The process of Claim 4, wherein the process action for training the neural network ensemble comprises an action of preparing each model image characterization from a model image depicting the face of a known person that it is desired to identify in the input image by,

extracting the portion of the model image depicting said face,

normalizing the extracted portion of the model image by resizing it to a prescribed scale if not already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area, and

cropping the extracted portion of the model image by eliminating unneeded portions of the image not specifically depicting part of the face of the subject to create a model face image.

6. The process of Clam 5, wherein the process action for training the neural network ensemble further comprises actions for:

categorizing the model face images by assigning each to one of a set of pose ranges into which its associated face pose falls;

for each pose range,

choosing a prescribed number of the model face images of each person being modeled which have been assigned to the selected pose range,

concatenating each of the chosen model face images to create a respective dimensional column vector (DCV) for each,

computing a covariance matrix from the DCVs,

calculating eigenvectors and corresponding eigenvalues from the covariance matrix,

ranking the eigenvalues in descending order,

identifying a prescribed number of the top eigenvalues,

using the eigenvectors corresponding to the identified eigenvalues to form the rows of a basis vector matrix (BVM) for the pose range; and

multiplying each DCV by each BVM to produce a set of PCA coefficient vectors for each model face image.

7. The process of Claim 6, wherein the neural network ensemble employed comprises a plurality of face recognition neural networks each of which has input and output units and is dedicated to a particular pose range, said face recognition neural networks output units being in communication with input units of a fusing neural network which has at least enough output units to allow a different output to represent each person it is desired to identify at each of the pose ranges, and wherein the process action of training the neural network ensemble further comprises the actions of:

for each face recognition neural network, inputting, one at a time, each of the PCA coefficient vectors associated with the pose range of the face recognition neural network into the inputs of the network until the outputs of the network stabilize;

initializing the fusing neural network for training;

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network associated the vector's particular pose range group until all the PCA coefficient vectors of every DCV have been input, and repeating until the outputs of the fusing neural network stabilize; and

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network

associated the vector's particular pose range group and assigning the active output of the fusing neural network as corresponding to the particular person and pose associated with the model image used to create the set of PCA coefficient vectors.

8. The process of Claim 7, wherein the process action of employing the neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person, comprises the actions of:

preparing the face region extracted from an input image by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing the extracted face region to the same prescribed scale if not already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area, and wherein the cropping comprises eliminated unneeded portions of the image not specifically depicting part of the face of the subject;

concatenating the prepared face region to create a DCV;

multiplying the DCV by each BVM to produce a set of PCA coefficient vectors for the extracted face region;

inputting each PCA coefficient vector in the set of PCA coefficient vectors into the respective face recognition neural network associated that vector's particular pose range group; and

identifying the active unit of the output of the fusing neural network and designating the person and pose previously assigned to that unit as the person and pose associated with the extracted face region.

9. The process of Claim 3, wherein the process action for designating the input image region to be an unknown person comprises the actions of:

training the neural network ensemble to identify the person associated with the characterized input image region to be an unknown person if

it does not match any of the model image characterization to a prescribed degree; and

employing the neural network ensemble to identify the person associated with the characterized input image region to be an unknown person if it does not match any of the model image characterization to the prescribed degree.

10. A face recognition system for identifying a person depicted in an input image, comprising:

a general purpose computing device; and

a computer program comprising program modules executable by the computing device, wherein the computing device is directed by the program modules of the computer program to,

capture model images, each of which depicts at least one person of known identity,

locate and extract regions within the model images, each of which depicts the face of a known person that it is desired to identify in the input image,

determine the face pose for each of the face regions extracted from the model images,

categorize each face region by assigning each to one of a set of pose ranges into which its associated face pose falls,

capture said input image,

locate and extract regions within the input images, each of which depicts the face of a person, and

for each face region extracted from the input image, comparing the face region to the face regions assigned to each pose range to identify the person depicted in the input image face region.

11. The system of Claim 10, wherein the program module for comparing comprises sub-modules for:

assessing the degree of similarity between the input image face region and the face regions assigned to each pose range;

ascertaining which face region assigned to the pose ranges exhibits the greatest degree of similarity to the face region extracted from the input image; and

designating the face region extracted from the input image to be the person depicted in the face region ascertained to exhibit the greatest degree of similarity to the input image face region.

12. The system of Claim 11, wherein the program module for comparing further comprises a sub-module for specifying that the face region extracted from the input image has a pose falling within the pose range to which the face region ascertained to exhibit the greatest degree of similarity to the input image face region was assigned.

13. The system of Claim 12, further comprising a program module for designating the face region extracted from the input image to be an unknown person if the degree of similarity between the input image face region and the face regions assigned to each pose range do not exceed a prescribed threshold.

14. The system of Claim 10, wherein the program module for comparing comprises sub-modules for:

training a neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person; and

employing the neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person.

15. The system of Claim 14, wherein the sub-module for training the neural network ensemble comprises a sub-module for preparing each face region extracted from said model images by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing each extracted face region to the same prescribed scale if not already at the prescribed scale and adjusting each region so that the eye locations of the depicted subject fall within the same prescribed area, and wherein said cropping comprises eliminating unneeded portions of the image not specifically depicting part of the face of the subject.

A2 cont

16. The system of Clam 15, wherein the sub-module for training the neural network ensemble further comprises sub-modules for:

(a) selecting a previously unselected one of the set of pose ranges;

(b) choosing a prescribed number of the prepared face images of each person being modeled which have been assigned to the selected pose range;

(c) concatenating each of the chosen prepared face images to create a respective dimensional column vector (DCV) for each;

(d) computing a covariance matrix from the DCVs;

(e) calculating eigenvectors and corresponding eigenvalues from the covariance matrix;

(f) ranking the eigenvalues in descending order;

(g) identifying a prescribed number of the top eigenvalues;

(h) using the eigenvectors corresponding to the identified eigenvalues to form the rows of a basis vector matrix (BVM) for the selected pose range;

(i) repeating actions (a) through (h) for each remaining pose range;

(j) multiplying each DCV by each BVM to produce a set of PCA coefficient vectors for each face image.

17. The system of Claim 16, wherein the neural network ensemble employed comprises a plurality of face recognition neural networks each of which has input and output units and is dedicated to a particular pose range, said face recognition neural networks output units being in communication with input units of a fusing neural network which has at least enough output units to allow a different output to represent each person it is desired to identify at each of the pose ranges, and wherein the sub-module for training the neural network ensemble further comprises sub-modules for:

for each face recognition neural network, inputting, one at a time, each of the PCA coefficient vectors associated with the pose range of the face recognition neural network into the inputs of the network until the outputs of the network stabilize;

initializing the fusing neural network for training;

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network associated the vector's particular pose range group until all the PCA coefficient vectors of every DCV have been input, and repeating until the outputs of the fusing neural network stabilize; and

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network associated the vector's particular pose range group and assigning the active output of the fusing neural network as corresponding to the particular person and pose associated with the model image used to create the set of PCA coefficient vectors.

18. The system of Claim 17, wherein the sub-module for employing the neural network ensemble to identify the person depicted in the input image face

region and the pose associated with the face of the identified person, comprises sub-modules for:

preparing the face region extracted from an input image by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing the extracted face region to the same prescribed scale if not already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area, and wherein the cropping comprises eliminated unneeded portions of the image not specifically depicting part of the face of the subject;

concatenating the prepared face region to create a DCV;

multiplying the DCV by each BVM to produce a set of PCA coefficient vectors for the extracted face region;

inputting each PCA coefficient vector in the set of PCA coefficient vectors into the respective face recognition neural network associated that vector's particular pose range group; and

identifying the active unit of the output of the fusing neural network and designating the person and pose previously assigned to that unit as the person and pose associated with the extracted face region.

19. The system of Claim 13, wherein the program module for designating the face region extracted from the input image to be an unknown person comprises sub-modules for:

training a neural network ensemble to identify the person associated with the input image face region to be an unknown person if it does not match any of the face regions assigned to each pose range to a prescribed degree; and

employing the neural network ensemble to identify the person associated with the input image face region to be an unknown person if it does not match any of the face regions assigned to each pose range to a prescribed degree.

20. A computer-readable memory for use in identifying a person depicted in an input image, comprising:

a computer-readable storage medium; and

a computer program comprising program modules stored in the storage medium, wherein the storage medium is so configured by the computer program that it causes a computer to,

input model images, each of which depicts at least one person of known identity,

locate and extract regions within the model images, each of which depicts the face of a known person that it is desired to identify in the input image,

determine the face pose for each of the face regions extracted from the model images,

categorize each face region by assigning each to one of a set of pose ranges into which its associated face pose falls,

capture said input image,

locate and extract regions within the input images, each of which depicts the face of a person, and

for each face region extracted from the input image, comparing the face region to the face regions assigned to each pose range to identify the person depicted in the input image face region.

21. The computer-readable memory of Claim 20, wherein the program module for comparing comprises sub-modules for:

assessing the degree of similarity between the input image face region and the face regions assigned to each pose range;

ascertaining which face region assigned to the pose ranges exhibits the greatest degree of similarity to the face region extracted from the input image; and

designating the face region extracted from the input image to be the person depicted in the face region ascertained to exhibit the greatest degree of similarity to the input image face region.

22. The computer-readable memory of Claim 21, wherein the program module for comparing further comprises a sub-module for specifying that the face region extracted from the input image has a pose falling within the pose range to which the face region ascertained to exhibit the greatest degree of similarity to the input image face region was assigned.

23. The computer-readable memory of Claim 20, further comprising a program module for designating the face region extracted from the input image to be an unknown person if the degree of similarity between the input image face region and the face regions assigned to each pose range do not exceed a prescribed threshold.

24. The computer-readable memory of Claim 20, wherein the program module for comparing comprises sub-modules for:

training a neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person; and

employing the neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person.

25. The computer-readable memory of Claim 24, wherein the sub-module for training the neural network ensemble comprises a sub-module for preparing each face region extracted from said model images by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing each extracted face region to the same prescribed scale if not already at the

prescribed scale and adjusting each region so that the eye locations of the depicted subject fall within the same prescribed area, and wherein said cropping comprises eliminating unneeded portions of the image not specifically depicting part of the face of the subject.

26. The computer-readable memory of Clam 25, wherein the sub-module for training the neural network ensemble further comprises sub-modules for:

(a) selecting a previously unselected one of the set of pose ranges;

(b) choosing a prescribed number of the prepared face images of each person being modeled which have been assigned to the selected pose range;

(c) concatenating each of the chosen prepared face images to create a respective dimensional column vector (DCV) for each;

(d) computing a covariance matrix from the DCVs;

(e) calculating eigenvectors and corresponding eigenvalues from the covariance matrix;

(f) ranking the eigenvalues in descending order;

(g) identifying a prescribed number of the top eigenvalues;

(h) using the eigenvectors corresponding to the identified eigenvalues to form the rows of a basis vector matrix (BVM) for the selected pose range;

(i) repeating actions (a) through (h) for each remaining pose range.

(j) multiplying each DCV by each BVM to produce a set of PCA coefficient vectors for each face image.

27. The computer-readable memory of Claim 26, wherein the neural network ensemble employed comprises a plurality of face recognition neural

networks each of which has input and output units and is dedicated to a particular pose range, said face recognition neural networks output units being in communication with input units of a fusing neural network which has at least enough output units to allow a different output to represent each person it is desired to identify at each of the pose ranges, and wherein the sub-module for training the neural network ensemble further comprises sub-modules for:

for each face recognition neural network, inputting, one at a time, each of the PCA coefficient vectors associated with the pose range of the face recognition neural network into the inputs of the network until the outputs of the network stabilize;

initializing the fusing neural network for training;

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network associated the vector's particular pose range group until all the PCA coefficient vectors of every DCV have been input, and repeating until the outputs of the fusing neural network stabilize; and

for each DCV, simultaneously inputting the PCA coefficient vectors generated from the DCV into the respective face recognition neural network associated the vector's particular pose range group and assigning the active output of the fusing neural network as corresponding to the particular person and pose associated with the model image used to create the set of PCA coefficient vectors.

28. The computer-readable memory of Claim 27, wherein the sub-module for employing the neural network ensemble to identify the person depicted in the input image face region and the pose associated with the face of the identified person, comprises sub-modules for:

preparing the face region extracted from an input image by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing the extracted face region to the same prescribed scale if not

already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area, and wherein the cropping comprises eliminated unneeded portions of the image not specifically depicting part of the face of the subject;

concatenating the prepared face region to create a DCV;

multiplying the DCV by each BVM to produce a set of PCA coefficient vectors for the extracted face region;

inputting each PCA coefficient vector in the set of PCA coefficient vectors into the respective face recognition neural network associated that vector's particular pose range group; and

identifying the active unit of the output of the fusing neural network and designating the person and pose previously assigned to that unit as the person and pose associated with the extracted face region.

29. The computer-readable memory of Claim 23, wherein the program module for designating the face region extracted from the input image to be an unknown person comprises sub-modules for:

training a neural network ensemble to identify the person associated with the input image face region to be an unknown person if it does not match any of the face regions assigned to each pose range to a prescribed degree; and

employing the neural network ensemble to identify the person associated with the input image face region to be an unknown person if it does not match any of the face regions assigned to each pose range to a prescribed degree.

30. The process of Claim 1, wherein the comparing, ascertaining, designating and specifying process actions comprises the actions of:

training a network ensemble to identify the person associated with the characterized input image region and their face pose, wherein the network

ensemble comprises, a first stage having a plurality of classifiers each of which is dedicated to a particular pose range and outputs a measure of the similarity indicative of the similarity between said characterized input image region and each of said model image characterizations associated with the particular pose range of the classifier, and a neural network as its second stage, said measure of similar of each classifier being in communication with inputs of the neural network which has at least enough outputs to allow a different output to represent each person it is desired to identify at each of the pose ranges; and

employing the network ensemble to identify the person associated with the characterized input image region.

31. The process of Claim 30, wherein the process action for training the network ensemble comprises an action of deriving each model image characterization from a set of model images of people, wherein each model image of the same person shows that person at a different face pose , said deriving action comprising:

extracting the portion of each model image depicting a face;

normalizing the extracted portion of each model image by resizing it to a prescribed scale if not already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area;

cropping the extracted portion of each model image by eliminating unneeded portions of the image not specifically depicting part of the face of the subject to create a model face image;

concatenating each of the model face images to create a respective model dimensional column vector (DCV) for each,

categorizing the model DVCs by assigning each to one of a set of pose ranges into which its associated face pose falls;

inputting the model DVC of the each model face image falling in a particular pose range, one at a time, to a pre-selected classifier dedicated to the particular pose range.

32. The process of Claim 31, wherein the process action of training the network ensemble further comprises the actions of:

initializing the fusing neural network for training;

simultaneously inputting the respective DVC of each model face image into all classifiers, until the DVC of every model image has been input, and repeating until the outputs of the neural network stabilize; and

simultaneously inputting the respective DVC of each model face image into all classifiers, and assigning the active output the neural network as corresponding to the particular person and pose associated with the model image used to create the DCV.

33. The process of Claim 32, wherein the process action of employing the network ensemble to identify the person depicted in the input image face region, comprises the actions of:

preparing the face region extracted from an input image by normalizing and cropping the extracted regions, wherein said normalizing comprises resizing the extracted face region to the same prescribed scale if not already at the prescribed scale and adjusting the region so that the eye locations of the depicted subject fall within a prescribed area, and wherein the cropping comprises eliminated unneeded portions of the image not specifically depicting part of the face of the subject;

concatenating the prepared face region to create a DCV;

inputting the DVC of the face region into all classifiers; and

identifying the active output of the neural network and designating the person previously assigned to that unit as the person associated with the extracted face region.

34. The process of Claim 33, further comprising a process action of specifying that the person designated as associated with the extracted face region has the face pose previously assigned to the identified active output.

35. A face recognition neural network ensemble for identifying a person depicted in an input image and a face pose range among a set of pose ranges into which the face of each identified person falls, comprising:

a plurality of face recognition neural networks each of which has input and output units and each of which is dedicated to a particular pose range; and

a fusing neural network whose inputs are in communication with the output units of said face recognition neural networks and which has at least enough output units to allow a different output to represent each person it is desired to identify at each of the pose ranges; and wherein

image feature characterizations derived from the face of a person it is desired to identify as depicted in the input image are respectively input into separate ones of the input units of the face recognition neural networks causing a single one of the output units of the fusing neural network to become active, thereby indicating the identity of the person whose face was depicted in the input image as well as the pose range associated with the pose of the depicted face.

36. The neural network ensemble of Claim 35, wherein the number of input units of each face recognition neural network equals the number of image feature characterizations derived from the face of the person to be identified as depicted in the input image.

37. The neural network ensemble of Claim 35, wherein the number of output units of each face recognition neural network at least equals the number of different people it is desired to identify in the input image.

38. The neural network ensemble of Claim 35, wherein the output units of each face recognition neural network output real values ranging from 0 to 1.

39. The neural network ensemble of Claim 35, wherein the number of input units of the fusing network is equal to the number of face recognition neural networks multiplied by the number of output units of any one of the face recognition neural networks, and wherein the number of output units of the fusing network is equal to the number of its input units.

40. The neural network ensemble of Claim 35, wherein the output units of each face recognition neural network are binary in that a particular output unit is active whenever it has the largest output amongst all the output units thereby representing a 1, and otherwise inactive thereby representing a 0.

41. The neural network ensemble of Claim 37, wherein there is one output unit of each face recognition neural network in addition to the number required to equal the number of different people it is desired to identify in the input image, and wherein at least one of the output units of the fusing neural network represents a person of unknown identity and unknown face pose.

42. A face recognition network ensemble for identifying a person depicted in an input image and a face pose range among a set of pose ranges into which the face of each identified person falls, comprising:

a plurality of classifiers each of which has input and output units and each of which is dedicated to a particular pose range; and

a fusing neural network whose inputs are in communication with the output units of said classifiers and which has at least enough output units to allow a different output to represent each person it is desired to identify at each of the pose ranges; and wherein

image feature characterizations derived from the face of a person it is desired to identify as depicted in the input image are respectively input into separate ones of the input units of the face recognition classifiers causing a single one of the output units of the fusing neural network to become active, thereby indicating the identity of the person whose face was depicted in the input image as well as the pose range associated with the pose of the depicted face.